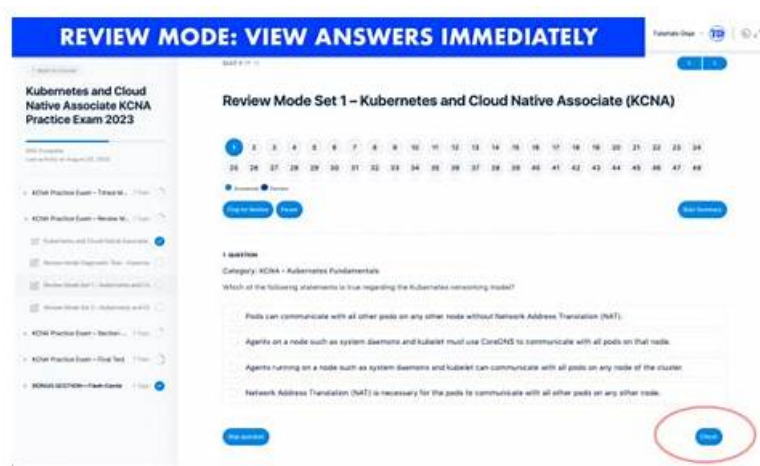


KCNA Valid Exam Topics & Valid KCNA Test Labs



BTW, DOWNLOAD part of Prep4away KCNA dumps from Cloud Storage: https://drive.google.com/open?id=1b0fRTS4k0ZfQ_YQtCusyFgrWPRvXG8v7

Our KCNA exam questions not only includes the examination process, but more importantly, the specific content of the exam. In previous years' examinations, the hit rate of KCNA learning quiz was far ahead in the industry. We know that if you really want to pass the exam, our study materials will definitely help you by improving your hit rate as a development priority. After using KCNA training prep, you will be more calm and it is inevitable that you will get a good result.

Linux Foundation KCNA (Kubernetes and Cloud Native Associate) Certification Exam is a globally recognized certification program designed for IT professionals who want to demonstrate their skills and knowledge in container orchestration and cloud-native technologies. KCNA Exam provides a comprehensive assessment of an individual's understanding of Kubernetes and other cloud-native technologies, including containerization, microservices, and networking.

>> **KCNA Valid Exam Topics** <<

Valid KCNA Test Labs, Simulations KCNA Pdf

There are three different versions of KCNA practice materials for you to choose, including the PDF version, the software version and the online version. You can choose the most suitable version for yourself according to your need. The online version of our KCNA exam prep has the function of supporting all web browsers. You just need to download any one web browser; you can use our KCNA test torrent. We believe that it will be very useful for you to save memory or bandwidth. In addition, if you use the online version of our KCNA Test Questions for the first time in an online state, you will have the opportunity to use our KCNA exam prep when you are in an offline state, it must be very helpful for you to learn in anytime and anywhere. If you think our products are useful for you, you can buy it online.

Linux Foundation is a non-profit organization that provides support and resources for open-source software projects. One of its most popular offerings is the Kubernetes and Cloud Native Associate (KCNA) certification exam. Kubernetes and Cloud Native Associate certification is designed for professionals who want to demonstrate their expertise in managing and deploying cloud-native applications using Kubernetes.

The KCNA Exam is a vendor-neutral certification, meaning that it is not tied to any specific cloud provider or technology. This makes it an attractive option for professionals who work with different cloud platforms and want to showcase their skills in a way that is recognized across the industry. KCNA exam covers a broad range of topics, including Kubernetes architecture, deployment, networking, and security.

Linux Foundation Kubernetes and Cloud Native Associate Sample Questions (Q17-Q22):

NEW QUESTION # 17

You're developing a new microservice for an application running on Kubernetes. The service requires access to a database hosted in a separate Kubernetes cluster. How would you ensure secure communication between your microservice and the database?

- A. Use a shared secret stored in an environment variable within the microservice pod
- B. Use an API gateway to proxy requests between the services.
- C. Configure a Kubernetes Ingress resource to route traffic to the database
- **D. Configure a service mesh to handle encryption and authentication between the services.**
- E. Set up a VPN connection between the two Kubernetes clusters-

Answer: D

Explanation:

A service mesh like Istio or Linkerd provides a layer of abstraction that can handle encryption, authentication, and authorization between microservices within and across Kubernetes clusters. It allows for secure communication without exposing sensitive credentials directly within the microservice code.

NEW QUESTION # 18

What is the main purpose of the Open Container Initiative (OCI)?

- A. Creating industry standards around container formats and runtimes for private purposes.
- **B. Creating open industry standards around container formats and runtimes.**
- C. Improving the security of standards around container formats and runtimes.
- D. Accelerating the adoption of containers and Kubernetes in the industry.

Answer: B

Explanation:

B is correct: the OCI's main purpose is to create open, vendor-neutral industry standards for container image formats and container runtimes. Standardization is critical in container orchestration because portability is a core promise: you should be able to build an image once and run it across different environments and runtimes without rewriting packaging or execution logic.

OCI defines (at a high level) two foundational specs:

* Image specification: how container images are packaged (layers, metadata, manifests).

* Runtime specification: how to run a container (filesystem setup, namespaces/cgroups behavior, lifecycle). These standards enable interoperability across tooling. For example, higher-level runtimes (like containerd or CRI-O) rely on OCI-compliant components (often runc or equivalents) to execute containers consistently.

Why the other options are not the best answer:

* A (accelerating adoption) might be an indirect outcome, but it's not the OCI's core charter.

* C is contradictory ("industry standards" but "for private purposes")-OCI is explicitly about open standards.

* D (improving security) can be helped by standardization and best practices, but OCI is not primarily a security standards body; its central function is format and runtime interoperability.

In Kubernetes specifically, OCI is part of the "plumbing" that makes runtimes replaceable. Kubernetes talks to runtimes via CRI; runtimes execute containers via OCI. This layering helps Kubernetes remain runtime-agnostic while still benefiting from consistent container behavior everywhere.

Therefore, the correct choice is B: OCI creates open standards around container formats and runtimes.

NEW QUESTION # 19

Which style of operations are preferred for Kubernetes and cloud-native applications?

- A. Imperative
- B. None of the above
- **C. Declarative**

Answer: C

Explanation:

<https://kubernetes.io/docs/tasks/manage-kubernetes-objects/declarative-config/#trade-offs>

NEW QUESTION # 20

Describe the different ways to manage persistent volumes in Kubernetes, including the concepts of static provisioning and dynamic provisioning. Provide examples for each approach.

- A. Static provisioning automatically generates PersistentVolumes (PVs) for each PersistentVolumeClaim (PVC) based on available resources. Dynamic provisioning manually defines PVCs but relies on Kubernetes to automatically select the appropriate StorageClass for the required storage type. This approach is not recommended for production environments but is useful for testing and prototyping.
- B. Static provisioning uses Kubernetes' built-in hostPath provisioner, which allocates persistent storage on the local node. In dynamic provisioning, administrators must create a StorageClass that links the PVC to a third-party cloud storage provider, such as AWS EBS. Static provisioning is best for development and testing environments where custom storage solutions are necessary, whereas dynamic provisioning is more suited for production applications where automated storage scaling is required.
- C. In static provisioning, the storage is pre-allocated, and PVs are manually bound to PVCs based on resource availability. In dynamic provisioning, Kubernetes allows automatic binding of PVs to PVCs without prior setup. This approach makes use of the default StorageClass or custom ones that support various storage backends such as NFS, Ceph, and cloud providers. Static provisioning requires more manual effort, while dynamic provisioning is easier to scale.
- **D. In static provisioning, you manually create PersistentVolumes (PVs) before deploying your application. This gives you more control over storage allocation, but it can be more complex for large deployments. In dynamic provisioning, you use a StorageClass to define storage characteristics, and the cluster automatically provisions PVs as needed. Dynamic provisioning simplifies storage management and allows for more scalable deployments. The YAML examples in option A demonstrate both approaches. The first example defines a static provisioned PV with a hostPath volume. The second example defines a dynamic provisioned StorageClass using the provisioner "kubernetes.io/gce-pd".**
- E. Static provisioning requires configuring both PVs and PVCs in YAML files, with explicit definitions for access modes and storage capacity. In dynamic provisioning, Kubernetes automatically generates PVCs as needed by the workloads based on predefined StorageClass settings. This reduces the need for manual intervention and simplifies large-scale deployments. The dynamic provisioning approach leverages persistent volumes provided by third-party cloud storage platforms such as Azure and Google Cloud.

Answer: D

Explanation:

In static provisioning, you manually create PersistentVolumes (PVs) before deploying your application. This gives you more control over storage allocation, but it can be more complex for large deployments. In dynamic provisioning, you use a StorageClass to define storage characteristics, and the cluster automatically provisions PVs as needed. Dynamic provisioning simplifies storage management and allows for more scalable deployments. The YAML examples in option A demonstrate both approaches. The first example defines a static provisioned PV with a hostPath volume. The second example defines a dynamic provisioned StorageClass using the provisioner "kubernetes.io/gce-pd".

NEW QUESTION # 21

How do you deploy a workload to Kubernetes without additional tools?

- A. Create a Bash script and run it on a worker node.
- B. Create a Python script and run it with kubectl.
- C. Create a Helm Chart and install it with helm.
- **D. Create a manifest and apply it with kubectl.**

Answer: D

Explanation:

The standard way to deploy workloads to Kubernetes using only built-in tooling is to create Kubernetes manifests (YAML/JSON definitions of API objects) and apply them with kubectl, so C is correct. Kubernetes is a declarative system: you describe the desired state of resources (e.g., a Deployment, Service, ConfigMap, Ingress) in a manifest file, then submit that desired state to the API server. Controllers reconcile the actual cluster state to match what you declared.

A manifest typically includes mandatory fields like apiVersion, kind, and metadata, and then a spec describing desired behavior. For example, a Deployment manifest declares replicas and the Pod template (containers, images, ports, probes, resources). Applying the manifest with kubectl apply -f <file> creates or updates the resources. kubectl apply is also designed to work well with iterative changes: you update the file, re-apply, and Kubernetes performs a controlled rollout based on controller logic.

Option B (Helm) is indeed a popular deployment tool, but Helm is explicitly an "additional tool" beyond kubectl and the Kubernetes API. The question asks "without additional tools," so Helm is excluded by definition. Option A (running Bash scripts on worker nodes) bypasses Kubernetes' desired-state control and is not how Kubernetes workload deployment is intended; it also breaks portability and operational safety. Option D is not a standard Kubernetes deployment mechanism; kubectl does not "run Python scripts" to deploy workloads (though scripts can automate kubectl, that's still not the primary mechanism).

